The Use of Metadata Visualisation to Assist Information Retrieval

Agata McCormac, Kathryn Parsons and Marcus Butavicius

Command, Control, Communications and Intelligence Division
Defence Science and Technology Organisation

DSTO-TR-2057

ABSTRACT

Metadata is descriptive information about data, which can be used to manage, locate or retrieve information. Although tabular presentations of metadata can be extremely useful, the exploitation of such information may be improved by visualisation. There are a number of information visualisation interfaces available, and many of these utilise metadata. However, on the whole, these approaches have not been objectively evaluated, and there is little information about their validity and reliability. This study analyses some of these techniques, highlighting their strengths and weaknesses, and the areas where further research is required. A study is proposed, in which a simple interface will empirically assess the efficacy of a metadata visualisation technique.

RELEASE LIMITATION

Approved for public release

APPROVED FOR PUBLIC RELEASE

The Use of Metadata Visualisation to Assist Information Retrieval

Executive Summary

Metadata is a useful tool and resource that is able to communicate large amounts of information. Metadata is used extensively in library catalogues, research environments, sales and marketing, financial analysis, manufacturing industries and in the areas of police and defence intelligence. Most commonly, metadata is presented in a tabular manner, with each metadata tag placed in a separate column. The use of such metadata tables has become common practice in many real-world applications.

In contrast, recent studies have demonstrated that users' ability to retrieve information can be significantly improved with the use of visualisations based on content similarity. Such techniques are rarely used in practical settings. For example, visualisation can provide an overview of the data; they allow quick and easy identification of clusters, trends, gaps or outliers; and enable a user to visually locate relationships and interactions in a way that is significantly easier than with metadata tables.

This report examines methods of metadata presentation and how best to integrate these with content-based visualisations. Details of metadata and information visualisation interfaces in the previous literature will be discussed as well as existing evaluations of these tools. Finally, a proposed study evaluating the effectiveness of (1) metadata tables, (2) content-based visualisation and a combination of (1) and (2) is proposed.

The literature survey revealed a lack of appropriate and robust research of the available tools. A major limitation of most studies is that they have been designed around small sample sizes, and this limits the ability to generalise results to other populations or other interfaces. Participants in these studies also have limited time to become accustomed to using a particular metadata interface, which means that results do not necessarily reflect potential performance.

Additionally, testing of more realistic tasks is required, since many experiments test very small data collections, and only examine very simple tasks. There is also a lack of appropriate and robust empirical evaluation, which means that the potential of metadata visualisation is unknown. We do not know if metadata interfaces are making appropriate use of our visual system, nor do we know how the metadata visualisation interfaces actually assist the user. Furthermore, it is still unclear why certain techniques are more effective than others.

Finally, this investigation was unable to locate any metadata visualisation tools that contain visualisations based on content similarity. The proposed study aims to address the limitations of previous research by using a simple interface to empirically assess whether the value of metadata is improved with the addition of a visualisation based on content similarity.

Authors

Agata McCormac
Command, Control, Communications and
Intelligence

Agata McCormac joined DSTO in 2006. She is a research scientist with the Information Access Group in C3ID where her work focuses on cognitive and perceptual psychology, information visualisation and interface design. She was awarded a Master of Psychology (Organisational and Human Factors) at the University of Adelaide in 2005.

Kathryn Parsons
Command, Control, Communications and
Intelligence

Kathryn Parsons is a research scientist with the Information Access Group in C3ID where her work focuses on cognitive and perceptual psychology, information visualisation and interface design. She obtained a Graduate Industry Linked Entrepreneurial Scheme (GILES) Scholarship in 2005, with LOD, where she was involved in human factors research, in the Human Sciences Discipline, specifically in the area of Infantry Situation Awareness. She completed a Master of Psychology (Organisational and Human Factors) at the University of Adelaide in 2005.

Marcus Butavicius
Command, Control, Communications and
Intelligence

Marcus Butavicius is a research scientist with the Information Access Group in C3ID. He joined LOD in 2001 where he investigated the role of simulation in training, theories of human reasoning and the analysis of biometric technologies. In 2002, he completed a PhD in Psychology at the University of Adelaide on mechanisms of visual object recognition. In 2003 he joined ISRD where his work focuses on data visualisation, decision-making and interface design. He is also a Visiting Research Fellow in the Psychology Department at the University of Adelaide.

Contents

1. Introduction

1.1 What is Metadata?

Metadata is, simply put, data about data. In other words, it is descriptive information about data, which can be used to explain, describe or locate an information resource, to make it easier to retrieve, use or to manage (NISO, 2004 and Tweedie, 1997). One of the most well-known uses of metadata is within a library catalogue. Catalogues are generally organised with metadata for each item within the library, providing information describing the author, the genre, the title, the publisher, the year it was published, any unique identifiers (such as the ISBN number), and the Dewey call number that would be used to locate the item. Any or all of these pieces of information can be used to search the catalogue. For example, a search for a specific author could quickly locate any other works by that person.

Music files also have metadata tags, in a format called ID3. This usually contains information such as the artist, the song title, the album title, the track length and the genre of music. Again, any of these pieces of information can be used to quickly search and locate specific tracks, to provide more information about the entire music collection, or to find similar or diverse tracks within the collection.

Metadata is traditionally divided into three main types, namely, descriptive metadata, structural metadata and administrative metadata. The examples of metadata described above are commonly referred to as descriptive metadata, as they can be used to describe a resource to enhance identification and retrieval (NISO, 2004). Structural metadata is, as suggested by the name, related to the organisation of the data such as the ordering of the sections that form a book (NISO, 2004). Finally, administrative metadata refers to the provision of information to assist the management of a resource, including when the data was created, the type of file, and other technical information (NISO, 2004).

This report will describe and examine a further example of metadata, which is based on the content similarity between documents. Essentially, utilising either human judgements or algorithms such as Latent Semantic Analysis (LSA), it is possible to obtain a measure of the similarity of every document in comparison to every other document. This is information about the resource, and is therefore metadata.

When used effectively, each document or piece of information has a number of metadata tags, and all documents or items sharing a particular metadata attribute should share the tags assigned to that attribute. For example, all documents regarding a particular topic should contain all of the same words or phrases used to describe that topic. This then allows knowledge representation, whereby any of the descriptive words could be used to assist in the retrieval of similar documents. Hence, metadata can reduce search time, as a user can simply explore using a metadata tag, which should then retrieve all of the relevant data regarding that specific attribute. In many cases, the metadata takes up less memory and contains less information than the main data and is therefore very useful to get the "gist" of a document or do a preliminary filter of information.

Beard and Sharma (1998) outline three main functions of metadata. It can be used to facilitate the overview of an information database, it can be used to enable a comparison between multiple pieces of information, and it can be used to provide a detailed description of individual items.

Despite these obvious benefits of metadata, the disadvantages must also be taken into account. Although many metadata tags comprise objective factual information, there are other metadata tags that are subjective (e.g. keywords assigned to a document), and hence, two people may not attach the same metadata tags to the same piece of information. Metadata can also be dependent on time and context, and words used to describe a piece of information at one point in time might differ considerably to the words used to describe that same piece of information in a different place, context, or in a different time period. Hence, it is important to consider these subjective limitations.

Also, most pieces of information could be allocated an extremely large number of metadata tags, and therefore metadata can become extremely extensive and complicated. For example, a photo could be described using information about the photo itself (referred to as the resource), such as the size of the photo, the camera used, the number of pixels etc. In addition, the photo could be described with contextual information, which not only includes the time and place that the photo was taken and who it was taken by, but also includes information about the scene in the photo, which could be described in great detail. Hence, a seemingly simple resource could quickly become quite complicated, and it may be difficult to put a threshold on the amount of information that should be provided in the metadata.

1.2 Metadata Visualisation

The section above has highlighted the potential value of metadata tags. However, metadata is usually presented in a tabular format, and when dealing with large amounts of information, long lists are not always intuitive for finding relevant documents. Humans have a highly developed visual system, and evidence suggests that people can understand the content of a picture far quicker than they can read and comprehend the meaning of text (Shneiderman, 1994). Therefore, in order to increase the value of metadata, information visualisations can be used, whereby the data is displayed in a more intuitive manner, which can assist users to quickly locate a relevant subset of the metadata, rather than requiring a comprehensive search through the traditional list (Weiss-Lijn, McDonnell & James, 2002).

Within the literature, there are a variety of terms used to describe the visualisation of metadata, and a variety of different methods used to visualise this information. Terms given to these interfaces include 'coordinated visualisations', 'dynamic queries', 'scientific visualisations' or 'visual information retrieval systems' (Shneiderman, 1994). These visualisations can range from simple interfaces, with one graphical display of the information, to extremely complicated coordinated visualisations, where several views of the data are provided in a technique known as 'Multiple Coordinated Views' (MCVs) (Grun, Gerken, Jetter, Konig & Reiterer, 2006).

Beard and Sharma (1998) claim that a single view of metadata is inadequate, and they highlight the importance of utilising multiple views, which allow for the dynamic nature of

the retrieval process. During the initial stages of information retrieval, a more overall picture is often required, then, after further examination of the data, the search criteria can usually be refined, and then more specific results are appropriate. Hence, the presentation, organisation and content of the metadata should change depending on the stage of the retrieval process (Beard and Sharma, 1998).

Most metadata visualisation techniques allow for the dynamic nature of information retrieval, and generally, the traditional list of metadata is presented together with a visual display of the information, often in a graphical form (Klein, Muller, Reiterer & Eibl, 2002). These interfaces allow the rapid and effective exploration of complex data sets, to discover, understand and explain the data (Dang, North & Shneiderman, 2001). Metadata visualisations can also reveal important information regarding clusters, trends, gaps or outliers in the data (Shneiderman, 1994), which is particularly important in situations involving complicated datasets, where correlations and relationships have not yet been found, or when a more open-ended exploration of the data is required (Dang et al., 2001). In a visual form these interactions can be quickly identified.

An example of this sort of visualisation is the layout of a street map. These maps are generally organised with the alphabetical list of street names and the location codes relating to each street on one side of the map, and the map itself on the other side, with location codes to find specific streets, and special symbols, to locate places of interest, such as schools, shops and other important buildings (Beagle, 1999). Both aspects of the map can be extremely useful, depending on the specific query. If the user is attempting to find a specific street, then the alphabetical list allows a quick directed search. However, the visualisation provides a more overall view of the area, which can also be extremely useful, particularly when a user is attempting to obtain more information about a particular suburb, as the map can quickly reveal information about the number and location of places of interest, which could be difficult to establish without the overview. In other words, it allows an examination of multiple sources of information simultaneously. An example may illustrate this point. If a user is interested in finding information about hotels in a particular area, then the alphabetical listing can provide the names and addresses of the relevant hotels. However, it is usually not until the map has been viewed, showing the location of the hotels in relation to other important places, that an informed decision can be made.

This highlights the different types of queries that are common when utilising an information-visualisation tool. Dang and colleagues (2001) report on an interface known as Dynamaps, which is used to present census data. This information visualisation tool provides users with a map and also consists of coordinated visualisations, dynamic-query sliders and other graphical representations of the data (see section 1.6 for more information on Dynamaps). This tool is usually used to answer specific questions such as "What is the population of my county?" or very general questions such as "Where is a nice place to live?" Although the specific population question could be answered quickly using the raw census data, the general question would be very difficult to answer effectively using the raw data alone. Hence, for this overall sort of question information-visualisation tools such as Dynamaps are extremely useful.

1.3 Common Attributes of Metadata Visualisations

As mentioned previously, there are a variety of different methods used to visualise metadata, and these different interfaces have a range of attributes. This section will outline a number of the common attributes used for metadata visualisations, which will be followed in a subsequent section by details of some specific metadata visualisation techniques or interfaces.

- *Search Function:* Arguably the most important feature of an information database is the search function, which can assist in the retrieval of information, by locating certain aspects of the data, based on user determined characteristics. Users can restrict their search to any aspect of the data, such as a certain subject area (Grun et al., 2006). Boolean searches include operator words such as "AND", "OR" and "NOT", and can also be used to assist in refining or extending the search. With most modern search functions, any of the metadata attributes, including title, author, year, media type, etc. can be used to limit a search.

- *Tables:* Tables are a basic but important attribute of most metadata visualisations, as they allow a huge amount of data to be displayed in a consistent manner (Grun et al., 2006). The data is arranged in rows and columns, and users can decide which attribute is most important for their purposes, and then often sort the data by that specific attribute (Grun et al., 2006).

- *Brushing and Linking:* Brushing and linking is another technique common to many coordinated visualisations. Essentially, when a piece (or pieces) of information is selected on one part of the interface, the equivalent piece of information is highlighted on the other parts of the interface. For example, if a piece of information is selected on a graphical display, that aspect of the data will then be automatically selected on the tabular display of information. This technique can assist the user to simultaneously obtain different sorts of information about the same piece of data. For example, the user can simultaneously see how the piece of information is related (or unrelated) to the rest of the dataset, and can also see the more detailed description of the particular piece of information.

- *Dynamic-Query Sliders:* A number of metadata visualisation techniques also utilise adjustable dynamic-query sliders, which is essentially a technique to filter the data. These sliders can be used to formulate queries, by altering them through a range of variables representing the attribute in question (Dang, et al., 2001). For example, with the census data, an adjustable dynamic-query slider could be provided for the metadata attribute, age. This could then be altered, to select a specific age which is of most interest to the particular query. As an attribute is selected, the visualisation is immediately updated, with the characteristic in question highlighted. When multiple dynamic query sliders are available, a further attribute can then be selected, to further narrow or filter the search. For example, with the census data, the attribute of gender could then be adjusted, to display only females. The visualisation would then have all females of the specific age group highlighted.

  Although the dynamic interfaces can be extremely useful, usability testing on tools utilising this technique have revealed that the system does not always respond

immediately, which can be frustrating and confusing for users (Rao & Mingay, 2001). Therefore, it is important that this possible problem is taken into account, and the system is designed to minimise lag-time. Also, usability testing of an interface with these sliders found that some participants had trouble locating the variable of interest on the slider (Rao & Mingay, 2001). However, in the interface in question the variables were not organised in any logical manner, and the authors suggest that this problem could be minimised by using an alphabetical listing of variables, or by allowing participants to select all of the variables that they would like to subsequently choose from using a drop-down menu (Rao & Mingay, 2001).

Li and North (2003) completed a usability study comparing the brushing technique with the dynamic query sliders. The study had a 'within subjects' design, with a counterbalanced order, which means that all participants performed tasks using both the brushing technique and the dynamic query sliders (Li & North, 2003). Results indicated that for complicated queries, including the comparison of data and the identification of trends, participants' performance was significantly faster with the brushing technique (Li & North, 2003). In contrast, participants could complete simple queries, such as tasks requiring the use of ranges, far quicker with the dynamic query sliders (Li & North, 2003).

- *Popup Windows:* Popup windows are often used in coordinated visualisations. These are similar to the brushing technique, whereby the user can select a piece of information in one section of the interface, and the details of this information will then 'popup' in a separate window or area of the interface. For example, if the interface is displaying documents, when one document is selected, the text of that document could then 'popup', allowing the user to see the piece of information in more detail. This is similar to the 'details' section of Microsoft Windows Explorer™ (Microsoft Corporation, 2001), in which a small version of the selected file is shown in the corner of the window.

  Some interfaces also have a popup of additional information which is shown when a user scrolls the mouse over a certain part of the display (Rao & Mingay, 2001). For instance, when participants place the mouse over a certain point in a scatterplot, the metadata for that item could then popup in a small window next to the mouse cursor. This can quickly give participants further information about the data, without over-complicating the display.

- *Graphical view:* Data can also be graphically displayed using methods such as the starfield display or scatterplot, in which the data is displayed with a collection of points, where the distance between the points demonstrates the relationship between those pieces of data. This view provides a more 'overall' picture of the data set, and can quickly reveal important information involving relationships, outliers, patterns or trends within the data, which can be very difficult to distinguish using only the traditional list.

  Traditionally, scatterplot displays present aspects of the metadata, with one metadata attribute along an x-axis, and another metadata attribute along the y-axis. For example, in Figure 4, based on Filmfinder, which is a tool designed to assist users to choose movies, the x-axis shows the year of production, and the y-axis displays the popularity of the films (Dang, et al., 2001).

In contrast, a more novel approach to the graphical display of data involves presenting the content similarity of documents. Essentially, the similarity of every document to every other document is quantified, and then these similarities are translated into a graphical display. The similarity judgements can be made via either human judgements, or algorithmic analysis. Hence, inter-object distance and spatial location is determined by the semantic similarity of the documents (Westerman, Collins & Cribbin, 2005). A study by Westerman and Cribbin (2000) assessed information retrieval using graphical displays, and found that information retrieval was more effective when the spatial mapping of the items was based on actual human ratings.

A study by Butavicius and Lee (submitted) empirically assessed four different visualisation techniques, based on human pairwise similarity judgements. The study required participants to retrieve information from documents that were represented by the points in the visualisations. The study found that the multidimensional scaling (MDS) visualisation had a significant advantage in terms of accuracy and the number of documents accessed. However, other usability testing has suggested that users are quite unfamiliar with the use of scatterplots, and, particularly when different displays are present and different variables can be selected, users often find it difficult to understand how the scatterplot relates to the information (Rao & Mingay, 2001). Consequently, it is important to ensure that clear instructions are given regarding the use of these graphical displays, and adequate training is then provided.

- *Additional components:* A number of metadata visualisations also include a zoom function and a pan or scroll function, which enables the user to focus in on a specific aspect of the interface. For example, if the visualisation includes a map, the user would be able to zoom in to a particular area of the map, to see it in more detail. Some interfaces also have a 'resize' function, where the users can decide which aspects of the tool are most useful for their purposes, and can then resize that tool, to put more focus on those important aspects.

1.4 Problems with Previous Studies Examining Metadata Visualisation

Regardless of how strong the theoretical background is, or how effective a visualisation technique may appear, if it fails to convey information effectively to users, then it is of little use (Kosara, Healey, Interrante, Laidlaw & Ware, 2003). Hence, it is extremely important that information visualisation techniques are created with the user in mind, and it is necessary to determine whether metadata visualisation interfaces actually assist the user. Despite the importance of testing the usability of these systems, there are a number of potential problems associated with examining metadata visualisation interfaces, and a number of limitations in previous studies of metadata visualisation tools. The limitations listed below will be expanded upon in the following section.

Table 2: Summary of Problems with Previous Studies Examining Metadata Visualisation

Limitation	Description
Lack of empirical evaluation	Research is rarely empirically based and interfaces are seldom assessed or compared to each other in an objective manner.
Inability to generalise findings	Studies generally relate to one interface, and due to differences between interfaces, findings can not be generalised.
Users' lack of familiarisation	Since metadata visualisation is unfamiliar, users' experimental performance is unlikely to represent potential performance.
Small sample sizes	Many studies have very small sample sizes, which limits the ability to generalise findings.
Influence of subjectivity	Some metadata have a subjective component and therefore definitional differences can influence the value of the metadata.
Limited results provided	Many studies only report the average score, which does not provide a true indication of performance for individual tasks.
Only simple tasks are assessed	Most of the experimental tasks are very simple with small datasets, so they are not representative of real world tasks.
Invalid conclusions	Studies often conclude that the interfaces are effective and useful, even if the data does not reflect this.

A review of the literature has highlighted that although there are many different metadata interfaces, there is very little empirical evaluation of these systems. There is also a lack of experimentation on display issues, such as the impact of colour, sound and other aspects of the interface (Shneiderman, 1994). Furthermore, different visualisation interfaces have many subtle differences in terms of design and features, and it is therefore very difficult to compare the results of different studies without some consistency in the way in which they are examined (Shneiderman, 1994).

Similarly, one of the main goals of user studies should be to determine *why* a technique is effective. Most user studies fail to address this, and instead, simply report whether or not their technique is effective, without much reference to the reasons behind the effectiveness (Kosara et al., 2003). Ideally, there should be more experimentation investigating these specific attributes, to determine the most user-friendly and performance enhancing designs.

Essentially, interfaces are generally assessed via usability testing rather than rigorous empirical analysis. Empirical testing involves a constrained experiment that follows strict guidelines and procedures, and usually involves an objective evaluation of multiple interfaces. In contrast, usability studies are generally only used to measure one tool, and involve far fewer users. Hence, it is more difficult to draw sound conclusions from usability testing.

Metadata visualisation is also an unfamiliar concept to most people. Therefore, a potential problem associated with experiments examining metadata visualisation is that people may have trouble becoming accustomed to the visualisation technique. In an experimental setting, it would be very difficult to give participants enough practice to reach the level of performance that could be expected once they become familiar with the system. Consequently, experimental results are likely to give an inaccurate indication of potential performance in the longer term with the metadata visualisation tool, and the results could

perhaps underestimate how well people could perform on the system once they are adequately trained. Similarly, evidence suggests that people can have trouble shifting between two or more completely different visualisations, and hence, it could be difficult for people to become accustomed to the technique (Limbrach, Muller, Klein, Ruiterer, Eibl, 2002). Although it is likely that people would eventually become accustomed to this, it is difficult to ascertain in the limited time users are normally studied for in experimental studies.

Furthermore, as mentioned previously, the studies examining metadata visualisation also tend to have small sample sizes. Experimenters generally maintain that participants must be from the correct population (Plaisant, 2004; Albertoni, Bertone & De Martino, 2005). For example, the people who would ultimately be using the system are the only people who can accurately test the system. Hence, if a metadata visualisation technique displays information for lawyers, experimenters would argue that only lawyers would be able to accurately judge the usefulness of the system. Since it is often difficult to get large sample sizes of these relevant populations, the studies are often carried out with an extremely small sample size, which means that it is difficult to determine whether the results are generalisable to the wider population.

Additionally, the effective use of information visualisation systems requires users to be intellectually engaged in the system, and it can be very difficult to achieve this in experimental conditions (Plaisant, 2004). The use of rewards for participation can improve commitment to the experiment, but performance may still be influenced. Hence, again it is possible that the results found in experiments may underestimate the ultimate performance that could be expected with a metadata visualisation technique.

A further problem that can arise when dealing with metadata is related to the vocabulary used, as the user must search with the same terms that were used in the initial categorisation of the data. Administrative metadata, such as the time or date, should remain objective, and hence, should not be affected. In contrast, descriptive metadata can have a subjective component, which means that if the assigned metadata tags were too specific, then it could be difficult for a user to determine the appropriate search term, and hence, the effectiveness and usefulness of the metadata visualisation tool could be limited (Albertoni et al., 2005).

Plaisant (2004) also suggests that the results of some studies could be biased, as they tend to present a summary of the results for all of the tasks, rather than presenting the individual results for each task. Since it is expected for some tasks to be completed more easily than others, this average score could easily result in a bias, where individual tasks with extremely high or low scores influence the overall score (Plaisant, 2004). Hence, in some situations, the scores for the individual tasks would give a far better indication of the performance of the interface.

Furthermore, evidence suggests that most empirical evaluations of visualisation systems involve only simple tasks, such as the location and identification of facts (Plaisant, 2004). More complicated tasks involving comparing, clustering and categorising are rarely covered. The interfaces tested also tend to lack the complexity that could be expected from real-life databases. For example, Li and North's (2003) empirical evaluation of the dynamic query sliders and the brushing technique used a limited dataset with only six attributes. They

concede that with a large and complex dataset there could be far more problems, particularly in relation to possible lag-time on the interface (Li & North, 2003).

Essentially, due to these problems associated with the experimental studies examining how metadata visualisations assist participants, there is little evidence to suggest an improvement in performance. Despite the fact that most studies do not find that metadata visualisation interfaces significantly increase performance, authors often claim that this lack of significant finding is not a reflection of limitations in the interface, but is rather an indication that participants are failing to make full use of the visualisation tool (Weiss-Lijn, McDonnell & James, 2001). Hence, despite the lack of evidence, authors often maintain that, if used properly, these visualisations will be more effective (Weiss-Lijn, McDonnell & James, 2001; Weiss-Lijn, McDonnell & James, 2002). Although it is possible that the optimal use of such interfaces could greatly assist performance, it is invalid to make this assumption. Unless usability testing or empirical evaluation can demonstrate it, a hypothetical or theoretical increase in performance is insufficient.

1.5 General Problems with Information Visualisation

In addition to the problems associated with studies examining metadata visualisations, there are also more general problems with the visualisation of information. The major problems and challenges that impact on the development of information visualisation have recently been summarised by Chen (2005). Chen (2005) identified the top ten unsolved information visualisation problems as belonging to three broad categories; user-centred issues, technical challenges, and issues that need to be addressed at a disciplinary level.

1.5.1 User-Centred Problems with Information Visualisation

Four user-centred issues have been identified and they include; usability, prior knowledge, understanding of elementary perceptual-cognitive tasks and education and training.

- *Usability*

Although research and growth of information visualisation has been rapid, as mentioned above, there has been a distinct lack of usability studies and empirical evaluations. On occasions where usability studies have been attempted, they are often conducted in an ad hoc manner and applied to particular systems, which limits the generalisabilty of findings. There is a clear need for the development of new evaluative methodologies that are able to address problems and needs that are specific to information visualisation (Chen, 2005).

- *Prior Knowledge*

The issue of prior knowledge is crucial to a user's ability to understand visualised information. Users require not only the knowledge of how to operate a system but they also require domain knowledge necessary for the interpretation of its content. This problem calls for the development of "adaptive visualization *[sic]* systems in response to accumulated knowledge of their users" (Chen, 2005, p.13). It highlights the importance of effective human-computer interaction (Johnson, 2004).

- *Understanding of Elementary Perceptual-Cognitive Tasks*

Research also needs to focus upon the level of discrepancy that exists between high level user tasks and the evaluation of the usefulness of various visualisation components. For example, tasks such as judging the relevance of information and browsing and searching through information "require a level of cognitive activities higher than that of identifying and decoding visualized objects" (Chen, 2005, p.13). A revision of elementary perceptual-cognitive tasks as they pertain to information visualisation is required.

- *Education and Training*

The problems associated with education and training can be overcome internally and externally (Chen, 2005). Internally researchers and practitioners need to share ideas, principles and skills and ensure that the language of information visualisation is consistent and comprehensible to all potential users. Externally there is a need for a public forum where members of the community can see the application of information visualisation and observe its contribution and potential (Chen, 2005).

1.5.2 Technical Problems with Information Visualisation

The technical challenges associated with information visualisation include the use of intrinsic quality metrics, scalability and aesthetics (Chen, 2005).

- *Intrinsic Quality Measures*

Intrinsic quality metrics are crucial for evaluation as they help to ensure quality. They are also used as a benchmarking tool. An example of such a metric is variance accounted for (VAF) equivalent to stress, used by the multidimensional scaling algorithms to measure the degree of fit to the final solution to the original data. It is the correlation between the original similarity matrix and the two dimensional similarities, calculated by the MDS algorithm in a particular solution. It has been shown that when VAF levels are low MDS solutions improve. Such quantifiable quality measures can facilitate the development and evaluation of algorithms and should provide the necessary evidence to strengthen advancements in information visualisation (Chen, 2005).

- *Scalability*

Scalability has consistently been a challenge for information visualisation. Primarily scalability involves developing methods to more effectively scale up computing algorithms. Although scalability will continue to be explored at a computing and hardware level, Chen (2005) suggests that in the future scalability issues also need to focus on the impact and influence of individual users.

- *Aesthetics*

The influence of aesthetics certainly cannot be overlooked. It is important to understand what representations make information visually appealing to users and what kinds of representations enhance insights. To achieve this Chen (2005) suggests that more empirical research needs to be conducted with the aim of discovering which visual properties make representations appealing to users. The majority of current research in "this area often focuses on graph-theoretical properties and rarely involves the semantics associated with the data" (Chen, 2005, p15). Researchers need to ensure that any changes to aesthetics are not only

visually appealing to users, but also translate to improved performance. This is important since research (Brath, Peters & Senior, 2005 and Purchase, 2000) has found that better aesthetics can be associated with poorer performance.

1.5.3 Problems at the Disciplinary Level

At the disciplinary level Chen (2005) identified three challenges; a paradigm shift from structure to dynamics, the issue of causality, visual inference and predictions, and finally the challenge of knowledge domain visualisation.

- *Paradigm shift from structures to dynamics*

A shift from the structure-centric paradigm to a dynamics paradigm will acknowledge the dynamic properties that underlie visualisation, and recognise that visualising changes over time (Chen, 2005). However, since change may not be rapid enough to attract attention, and since most visualisations lack trend detection systems, Chen (2005) also emphasises the importance of "interdisciplinary collaborations between the data mining and artificial intelligence communities" (Chen, 2005, p15).

- *Causality, visual inference, and predictions*

Information visualisation is a powerful tool that can enable users to find causality, make visual inferences and test predictions and hypotheses, and therefore "users need to freely interact with raw data as well as its visualizations" (Chen, 2005, p15). According to Chen (2005), this discovery process could be greatly enhanced through the use of multiple coordinated views. The successful achievement of this goal involves the development of algorithms that can filter out noise and process conflicting evidence and information.

- *Knowledge domain visualization*

Chen's (2005) final information visualisation challenge is holistic, incorporating aspects of all nine problems. This problem takes into account social construction, highlighting the fact that whilst, on the one hand, information is relatively stable, on the other hand, knowledge must consider the value or relevance of the piece of information. This challenge is large in scale, scope and duration and, if solved successfully, can potentially be applied to a wide range of subject areas (Chen, 2005).

1.6 Specific Information Visualisation Techniques

There are a vast number of different information visualisation interfaces available, either commercially, or for research purposes. Due to the large number of different information visualisation techniques, this paper will provide an overview of only a sample of the available interfaces. The following section will describe the interfaces, focusing on the strengths and weaknesses, and the empirical evidence of their effectiveness.

1.6.1 University of Maryland

The University of Maryland's Human-Computer Interaction Laboratory have a large research area focusing on information visualisation, and they have been involved in the development

of a number of user interfaces. This section will highlight some of the relevant information visualisation interfaces, highlighting the features, strengths and weaknesses of the tools.

- *Spotfire*

Spotfire is a commercially available tool that is used in over 1000 organisations around the world (Spotfire, 2006). It is a decision-making tool that is able to display trends and patterns, locate outliers and identify unexpected relationships in data (Spotfire, 2006). The Spotfire tool uses a range of interactive information visualisation techniques including a starfield display. A starfield is essentially "an interactive scatterplot with additional features for zooming, filtering, panning, details-on demand etc." (Ahlberg, 1996, p26). Spotfire provides an example of a visualisation system where users are able to create visualisations, manipulate objects within the visualisation, and complete high level exploration tasks (Ahlberg, 1996).
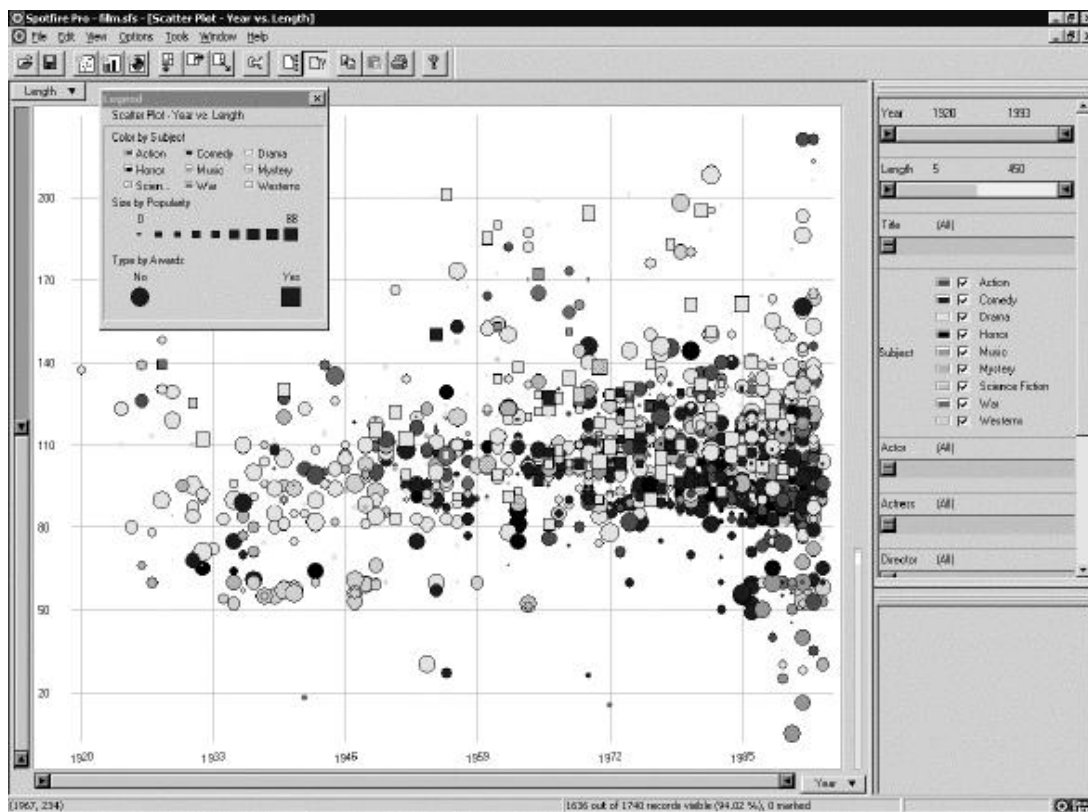
Figure 1: Spotfire (Dang et al., 2001)

Spotfire has been applied to a wide range of areas including; clinical settings, research environments, financial analysis, sales and marketing, the manufacturing industry and intelligence and defence activities (Spotfire, 2006). Despite the fact that this tool is used by various industries, there is a distinct lack of research into its effectiveness. Spotfire has not been vigorously assessed for its usability or validity nor has it been compared to other visualisation tools and interfaces. A screenshot of the Spotfire interface can be seen in Figure 1.

- *Snap-Together Visualizations*

Snap-Together Visualization (STV) allows users to explore complex information without computer programming skills. It is unique because it allows users to quickly construct their own coordinated visualisation interfaces that are specific to their customised data. A coordinated visualisation interface consists of a set of visualisations, which can interact, portraying the relationship that exists between them. A variety of methods can be used to explore the metadata and these co-ordinations include; brushing and linking, overview and detail, drill down and synchronized scrolling (North & Shneiderman, 2000a, 2000b, 1999b). For example, the census data displayed in Figure 2 shows a number of different views of the data, with '*Montgomery, MD*' highlighted in each section. This therefore shows multiple aspects of that particular region.
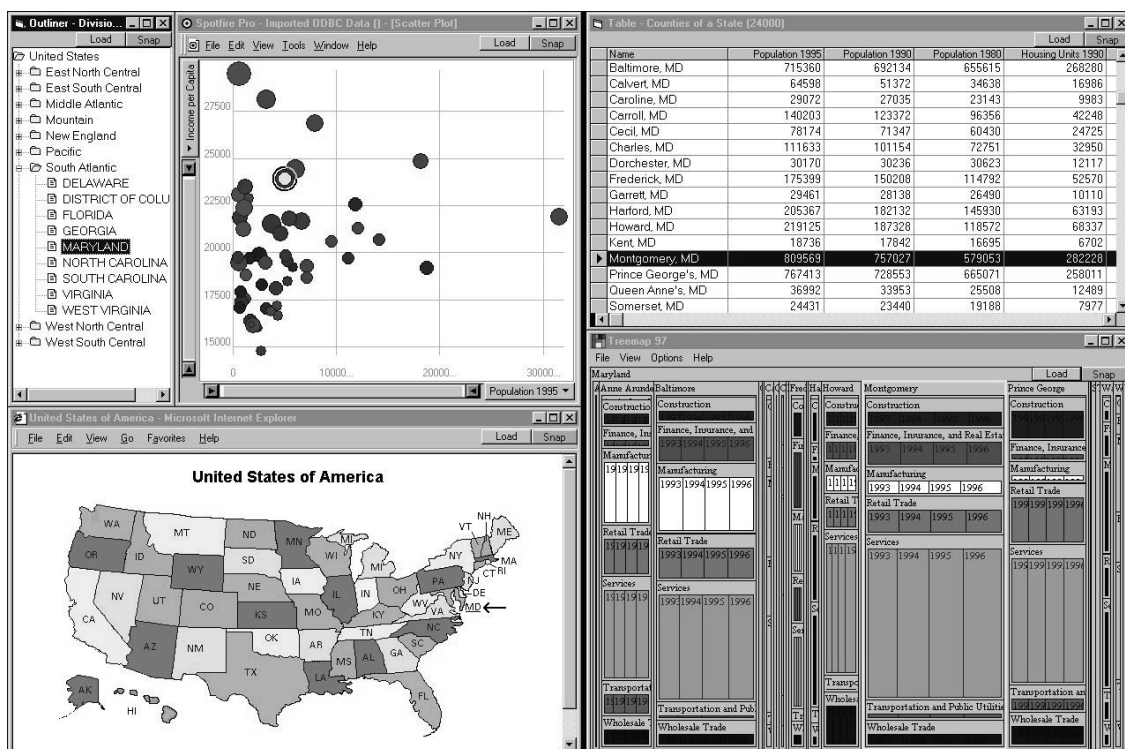
Figure 2: SnapTogether Visualisations

The visualisations and co-ordinations that a user may need will vary for each situation and will depend upon the features and structure of the data, the task and outcomes of a user, and it will also depend upon individual differences associated with a user, such as their experience, prior knowledge and user preference (North & Shneiderman, 2000a, 2000b). The strength of STV lies in its ability to communicate with other independent visualisation tools, through the use of an application programming interface. STV is able to rapidly construct coordinated visualisation interfaces to explore and navigate data and its relationships (North & Shneiderman, 2000a, 2000b). The STV interface has been applied to a wide range of data and information including; Census Bureau data, photo libraries, web logs, mailing lists, technical report databases and case law databases (North & Shneiderman, 2000b).

Two studies were conducted to evaluate the usability of the interface. The first study aimed to measure how successfully users were able to coordinate their own visualisation interfaces. The study found that overall subjects were able to quickly understand and use coordinated views and subjects were able to construct their own coordinated interfaces. Results from the second study were also encouraging (North & Shneiderman, 2000a, 2000b, 1999b). The second study compared the benefits of using views coordinated with STV to those coordinated with independent views or single views (North & Shneiderman, 2000a, 2000b, 1999b). Subjects performed nine different browsing tasks over the three different interfaces. The nine tasks ranged from easy to difficult. It was determined that on average, participants were significantly quicker when using the Snapped-Views, with participants approximately 80% quicker for easy tasks, and 30-50% quicker when completing difficult tasks (North & Shneiderman, 2000b).

Although the results of both studies are quite promising and positive they need to be treated with a degree of caution. The studies consisted of small sample sizes with six subjects participating in the first study and eighteen subjects participating in the second study. No further vigorous research or testing has been conducted to either replicate the findings or to further test the usability and validity of the tool.

- *Dynamaps and Census data interfaces*

Dynamap is a CD-ROM based tool that has been used by Census Bureau staff in America to facilitate viewing and analysis of map related data and information (Roa & Mingay, 2001). Dynamap provides an example of a real-world application of Snap Together Visualization. The interface was developed by The University of Maryland's Human Computer Interaction laboratory. Figure 3 shows an example of the interface, with brushing between a scatterplot and map revealing that the high income, highly educated states are in the northeast (Dang et al., 2001).
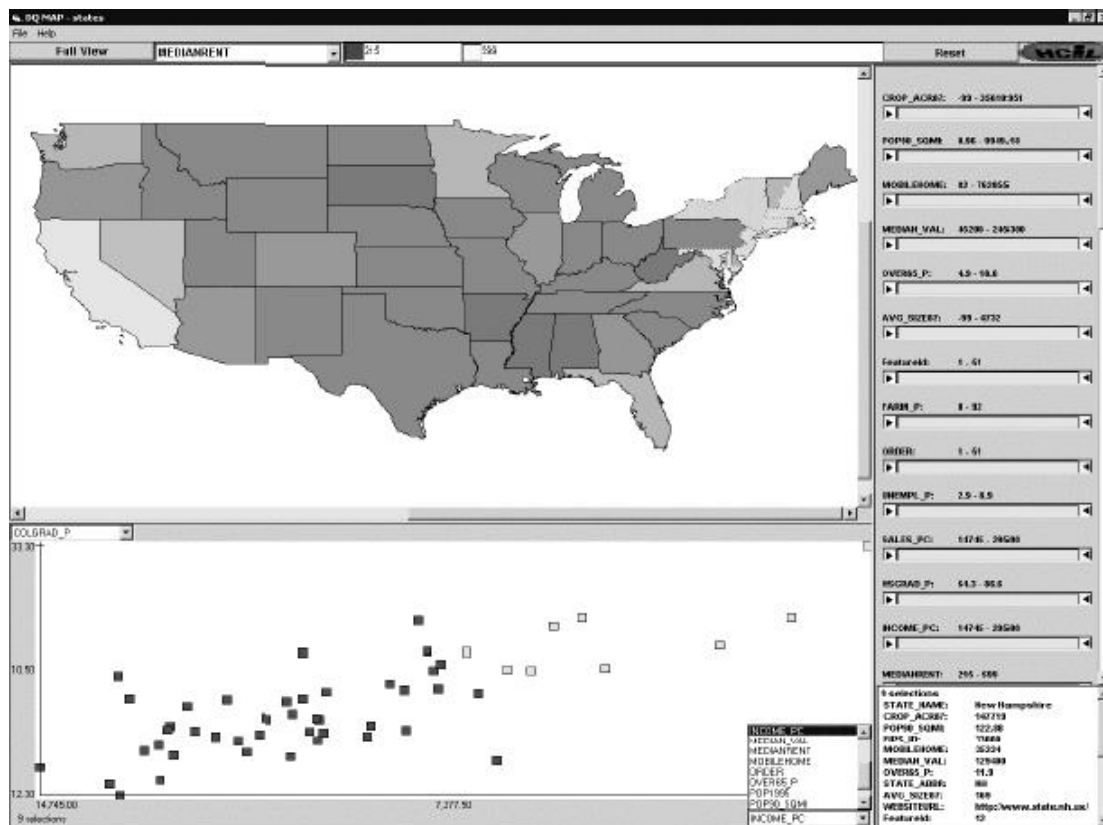
Figure 3: Dynamaps (Dang et al., 2001)

The University also conducted usability testing and made recommendations for future improvements (Roa & Mingay, 2001). Some of the problems that were identified included time lag issues, problems with the design of zoom controls and confusing and unclear instructions. It was also not clear to users what information was selected or deselected. Following the completion of usability testing, recommendations for improvements to overcome these challenges were made. However, no further research has been completed to measure the reliability and validity of the recommendations and suggested changes. Although users found the interface needed to be more intuitive and more easy-to-learn, they also found the tool to be functional and user rich (Roa & Mingay, 2001). However, once again further research and testing needs to be conducted to measure the usability and effectiveness of the interface.

- *FilmFinder*

The FilmFinder is a tool designed to assist users to choose movies. As displayed in Figure 4, it consists of a series of dynamic query sliders and a starfield display, with the x-axis representing time, and the y-axis representing popularity (Ahlberg & Shneirderman, 1994). The interface also includes dynamic query sliders representing title, actor, actress, director, length and rating, and the technique includes a zoom feature, so users can focus on a particular time or aspect of the popularity scale (Ahlberg & Shneirderman, 1994). The different genres (including drama, mystery, comedy, western, horror, action etc) are displayed in different colours, and it is possible to display only specified genres. When there are fewer than twenty-five movies, the titles are automatically displayed, and when users click on one of

the titles, more details of the title pop up in the display (Ahlberg & Shneirderman, 1994). For example, users could select Harrison Ford in the actor slider, and they could specify that they would like to see a movie with a G or PG rating. Then, all movies fitting those criteria would be displayed on the interface, allowing participants to make a more informed decision about their movie choice.
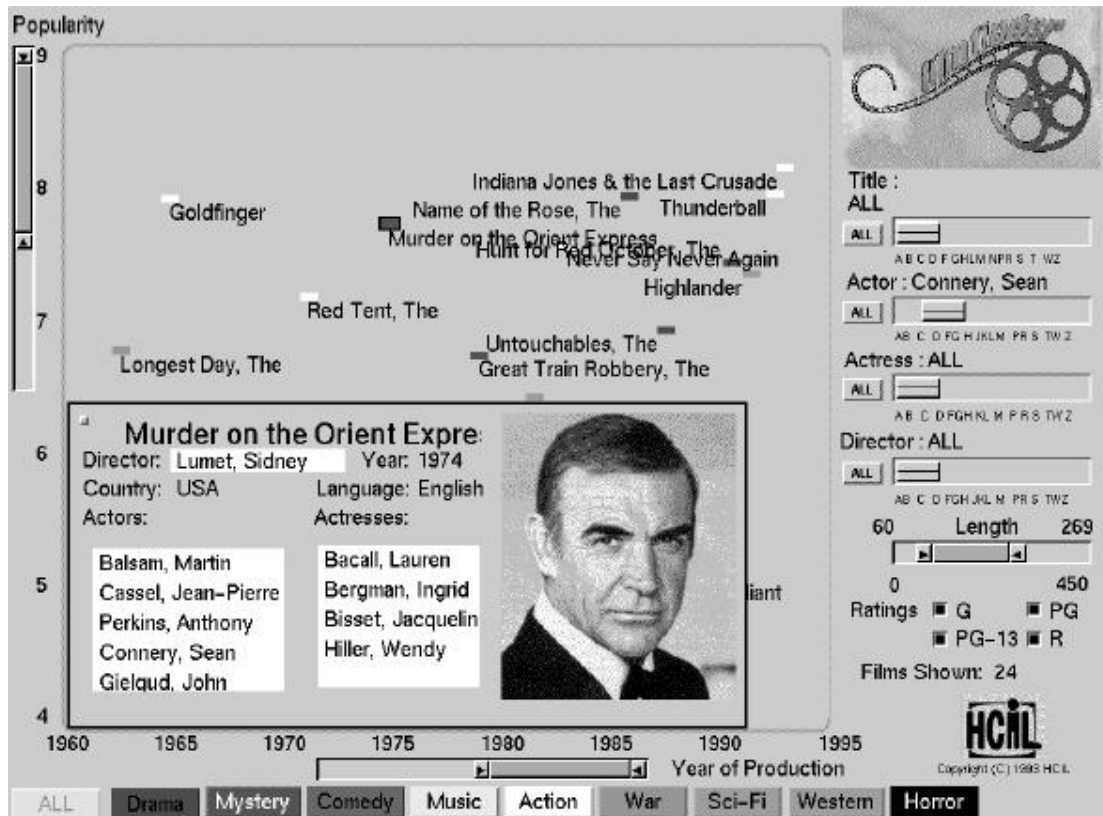

Figure 4: FilmFinder (Dang et al., 2001)

Unfortunately, like many of the other visualisation interfaces, there appears to be a lack of empirical testing of this technique. Also, the FilmFinder could be further improved, by providing a 'fuzzy' searching capacity, where users could search for films that are similar to a specified title (Ahlberg & Shneirderman, 1994). It is also important to test and enhance the usability of the interface, particularly in regards to the lag time (Ahlberg & Shneiderman, 1994).

- *HomeFinder*

The HomeFinder is a visualisation technique that provides real estate information for potential home buyers. The interface includes a map of the geographic area and a number of dynamic query sliders, allowing users to select homes according to characteristics such as location, cost, number of bedrooms and other features (including fireplace, garage, new construction and central air conditioning) (Shneiderman, 1994). Users can adjust the sliders to select certain features and a particular price range. Distinct markers (such as their place of work) can also be shown, and users can choose an acceptable distance from these markers, which can be displayed on the interface (Ahlberg & Shneiderman, 1994). For example,

16

Figure 5 provides an example where the user has specified that they would like a home in a price range between $50k and $500k, with between two and four bedrooms, which is less than 19 kilometres from 'A', and less than six kilometres from 'B'. The homes that meet those criteria are then shown on the starfield display.

Figure 5: HomeFinder (Dang et al., 2001)

An empirical study of this interface utilised real estate data and compared the participants' performance when using the HomeFinder to their performance when using a paper listing of the data and a natural-language version (Shneiderman, 1994). The study found that response times were significantly quicker using the HomeFinder, and subjective ratings appeared positive towards the interface (Shneiderman, 1994). However, the study was completed with only eighteen participants, and detailed results of the study were not provided. It is therefore difficult to make accurate conclusions regarding the usability or usefulness of the tool.

1.6.2 University of Konstanz

The University of Konstanz also have a human-computer interaction group, and they have researched and developed a number of relevant information visualisation interfaces. In this section, a number of these interfaces will be described, providing details of some of the features and strengths and weaknesses of the techniques for the visualisation of metadata.

- *SuperTable + Scatterplot visualisation*

The 'SuperTable + Scatterplot' visualisation was designed to be as user friendly as possible (Limbach, Muller, Klein, Reiterer, & Eibl, 2002). Rather than providing a number of different

attributes that the user can choose from, the SuperTable + Scatterplot combines these different features into the one table, with the intention that the users will feel like they are using one visualisation in different states, rather than multiple visualisations (Klein, Muller, Reiterer, & Eibl, 2002). Theoretically speaking, this technique should help to minimise problems associated with people having trouble shifting between two or more different visualisations (Limbrach et al, 2002).

The SuperTable can be altered through different levels, to provide more or less information to users. In the first level, the height of the rows is extremely small, so that users are provided with an overview of the entire document collection (Limbach et al., 2002). At this level, the rows are too small to permit text, so colours and bars are used to represent different search terms. As the user moves through the subsequent levels, more specific information is provided for each document. For example, in Level 3, the title, URL and abstract of documents is presented, whereas in the previous levels, there was not enough room to show this. In the final level of the SuperTable, users are presented with the document itself (Klein, Muller, Reiterer, & Limbach, 2003).

The idea behind this technique is that users can change the features shown with as much ease as possible. Although a number of visualisation interfaces have been designed using variations on this technique, (for example INSYDER, INVISIP and the Visual Metadata Browser (VisMeB) displayed in Figure 6) (Albertoni, Bertone, & De Martino, 2005), again, there is little empirically sound research justifying the advantages of this technique.
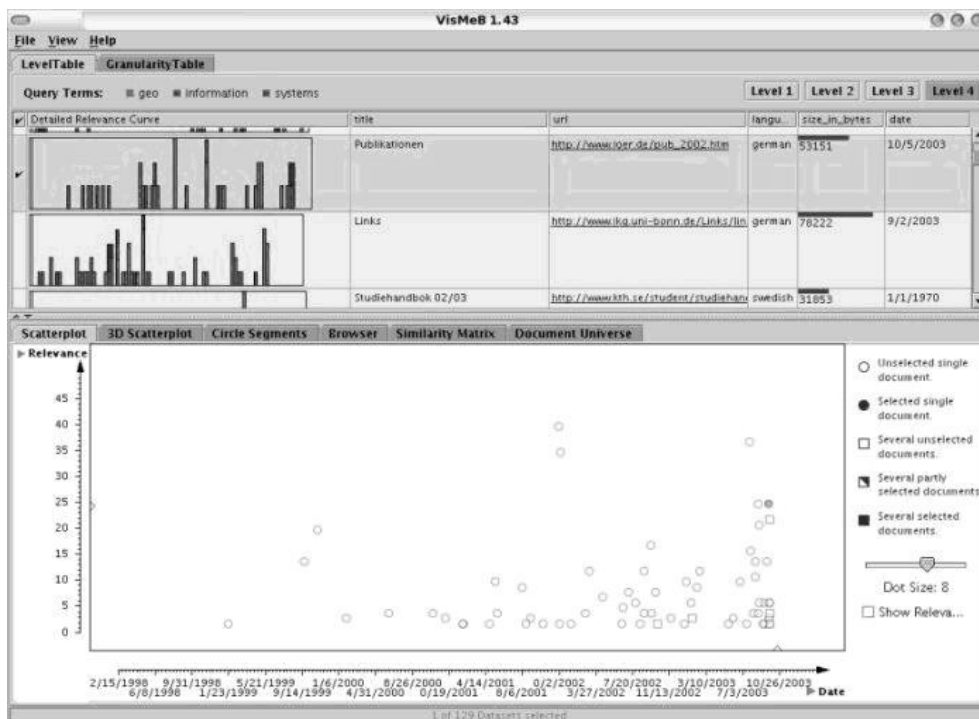
Figure 6: VisMeB (an application of the SuperTable + Scatterplot visualisation)

- *MedioVis*

MedioVis is a visual interface for searching and exploring multimedia libraries, catalogues and databases (Grun et al., 2005). It is a user-centred tool, which consists of multiple coordinated views, and was designed to assist non-expert users in the search and exploration of information databases (Grun et al., 2005). Grun and colleagues (2005) highlight the importance of aesthetics in the design of the MedioVis interface, suggesting that usability from an ergonomic perspective does not necessarily correlate with how a user would view the usability of an interface. Hence, an effort was made to enhance the attractiveness of the interface, and it offers visually appealing views of the data, such as multimedia items, including pictures, sound files and media clips. However, as mentioned previously, enhanced aesthetics does not always translate to an increase in performance (Purchase, 2000), and therefore it is important to evaluate the effectiveness of the interface.

Figure 6: MedioVis

The interface also includes a search function, a tabular display of the information, with columns for the different metadata attributes, a pop-up display of more detailed information, which is shown when a particular piece of data is selected, a zoom function, and a scatterplot-like visualisation, referred to as a 'graphical view'. A study compared the objective efficiency and subjective usability of MedioVis and KOALA, which is a web-based library catalogue system (Grun, et al., 2005). Twenty-four subjects took part in the experiment, and a counter-balanced design was used, which means that half of the participants completed the tasks on MedioVis first, and half of the participants completed the tasks on KOALA first (Grun et al., 2005). Results indicated that participants were able to complete tasks significantly faster with

MedioVis, and MedioVis was also favoured in the qualitative findings. However, the version of MedioVis used in the evaluation did not include the zoom function or the graphical view, and more empirical evaluations of the system are therefore required.

1.7 The Current Study

As indicated in this report, there is a lack of empirical research testing the many metadata visualisation interfaces that have been developed. There are also a number of deficiencies in many of the studies that have been completed, and hence, it is difficult to determine the effectiveness of the interfaces currently available. Furthermore, the metadata interfaces currently available tend to involve a graphical representation of the relationship between two metadata attributes. For example, the *FilmFinder* visualisation shown in Figure 4 provides a graphical representation of the relationship between year of production (on the x-axis) and popularity (on the y-axis).

In contrast, the proposed experiment is unique because it involves the graphical representation of a single metadata attribute, namely, content similarity. In this visualisation each piece of data will be represented by a point and the distance between these points indicates the relationship between the pieces of data. Such a visualisation is important, since previous studies (Butavicius & Lee, submitted) have indicated that information retrieval can be assisted with visualisations based on content similarity.

However, these displays only show content similarity, and do not provide useful metadata such as the title, author and year of a document. In many real-life applications, metadata is currently used as the predominant method of information access, and document similarity is rarely used.

Therefore, it is important to test whether the addition of a new capability (i.e., content similarity) impedes or improves upon the search strategies currently in operation (i.e., metadata search). In the literature there are no studies that have measured the performance of graphical displays based on content similarity and tabular presentations of metadata. The proposed study will assess this deficiency, empirically examining whether content similarity adds to metadata presented in a tabular display.

2. Proposed Methodology

2.1 Participants

The sample will consist of approximately 50 university students. It is anticipated that students will be recruited from third year level or higher. This is important as it should ensure that the participants represent the anticipated customer employees, in relation to age, gender and academic qualifications.

2.2 Apparatus and Measures

2.2.1 Demographic Questionnaire

Participants will be asked to complete a demographic questionnaire. This will ask questions regarding their age, gender, education level, area of study, and other research experience. The demographic questionnaire may also contain questions aimed at determining their previous experience with using metadata interfaces.

2.2.2 Cognitive Abilities Test

Participants will also be required to complete a short intelligence test or test of cognitive abilities. Research suggests that judgements of topicality may be influenced by these factors (Janes, 1994) and it is therefore important to obtain a measure of intelligence to determine whether cognitive abilities affect performance.

Also, since the completion of this experiment will require comprehension skills, a test of English comprehension or vocabulary will also be used, which should provide an indication of whether participants had difficulties with understanding the experiment.

2.3 Design and Procedure

Research assistants will be provided with detailed instructions to follow for carrying out the experimentation, which will ensure that all participants receive the same information. Before the experiment begins, participants will be given an information sheet and consent form, explaining their participation in the study. They will then be asked to complete a demographic questionnaire, followed by a short test of cognitive ability.

The study will consist of three conditions:
- An interface with a traditional list of metadata;
- An interface with metadata presented in a scatterplot type visualisation; and
- An interface with a coordinated visualisation, showing both the traditional list of metadata and the graphical view.

A program will be developed, which consists of the following attributes that were described in more detail in Section 1.3. The interface will consist of:

- A graphical view, which is a scatterplot visualisation, where the inter-object distance and spatial location is constructed using human pairwise similarity judgements. In other words, the content similarity of the documents will determine their location within the graphical display;
- A table, which is essentially a traditional tabular list of the metadata;
- A pop-up window, which will show the text of a selected document.

When participants are utilising the coordinated visualisation, the scatterplot and table will be linked, which means that a piece of information selected in one part of the interface will be highlighted in the other parts of the interface. For example, if a document is selected from the tabular view of metadata, the corresponding point will be highlighted in the scatterplot visualisation, and the text of the document will be shown in the pop-up window. An example of the interface is shown in Figure 7.

The data used will be newspaper articles, from the TREC 8 document collection. For each article, the metadata will include:
- Author,
- Source (the newspaper that the article is from),
- Date,
- Length,
- Publisher,
- Section (the section of the newspaper).

When utilising the tabular display or the coordinated visualisation participants will be able to sort the columns in ascending or descending order according to any one of the attributes. In this way they can prioritise the documents according to metadata attributes or examine variables in the entire document set across an attribute.

Participants will be given a research topic and will be asked to find all the documents that they would use to complete a research report on that topic. A repeated measures design will be used, which means that all participants will complete the experiment with each of the visualisation techniques. Hence, all participants will use the traditional list of metadata, the scatterplot visualisation of the documents, and the coordinated visualisation with both the traditional list of metadata and the graphical view.

The repeated measures design is used as it increases statistical power and reduces within group variance (this is because the same participants are in each of the conditions). To control for possible learning effects or fatigue, all conditions will be completed in a counterbalanced order, and three different sets of documents will be randomised across the conditions.

Each document set will consist of 60 documents, and 50% or 30 of those documents will be irrelevant to act as noise. This is important because it provides a better representation of a 'real' research task, where there may be a great deal of irrelevant information. The other 30 documents will be from one of three different research topics. For each condition participants will be asked to complete three questions, each of which will require the selection of documents that would be used to write a report on the specified topic.

The program used will not only record the documents that participants select, it will also record all other user interaction with the interface, including the documents viewed, the order in which documents are viewed, and all timing information. It is hypothesised that document selection will be easiest with the coordinated visualisation.

It is anticipated that participants should be able to complete each question in approximately 10 minutes. Hence, the experiment should take less than two hours, including the completion of the demographic questionnaire and cognitive abilities test. It will be necessary to complete a small pilot study to verify this length of time, and to reveal any other concerns.

Figure 7: An example of the interface

3. Summary

An information resource can be described with metadata tags, which can greatly assist in managing or locating data. Information visualisation is argued to be a more intuitive manner of displaying data, and can enhance the ability to comprehend information.

The hypothesis of this report is that metadata can be further improved by combining its usefulness with the many advantages associated with information visualisation. Essentially, lists of metadata attributes are not always very intuitive, whereas a visual display of information can utilise humans' highly developed visual systems, which can increase understanding and interpretation.

Many interfaces and systems comprising metadata visualisation have been developed, including *Spotfire* (Spotfire, 2006), *FilmFinder* (Ahlberg & Schneirderman, 1994), *MedioVis* (Grun, et al., 2005) and the *SnapTogether Visualisation* (North & Schneirderman, 2000a, 2000b). However, there is a distinct lack of empirical evaluation of these techniques.

The existing research tends to be limited by small sample sizes and poor statistical reporting. Furthermore, since metadata visualisations are unfamiliar to most people, it is also possible that the potential performance could be higher than the performance seen in experimental conditions.

By highlighting the problems associated with the previous studies examining metadata visualisation, this study has emphasised the need for more research in this area. The proposed study uses a simple interface to ascertain whether performance in an information retrieval task can be improved with the visualisation of content similarity, which is a metadata attribute. More specifically, it is addressing the question of how content based similarity displays add or detract from conventional metadata displays using tabular format.

4. References

Ahlberg, C (1996). Spotfire: An Information Exploration Environment. *SIGMOD Record, Vol.25, No.4*, 25-29.

Ahlberg, C. & Shneiderman, B. (1994). Visual Information Seeking: Tight coupling of dynamic query filters with starfield displays. *Proceedings of the 1994 ACM conference on Human Factors in Computing Systems (CHI)*, New York, New York.

Albertoni, R., Bertone, A., & De Martino, M. (2005). Visualization and semantic analysis of geographic metadata. *Proceedings of the 2005 Workshop on Geographic Information Retrieval*, Bremen, Germany.

Beagle, D. (1999). Visualization of metadata. *Information Technology and Libraries, 18(4)*, 192-199.

Beard, M.K., & Sharma, V. (1998). Multilevel and Graphical Views of Metadata, *Proceedings of the Fifth International Forum on Research and Technology Advances in Digital Libraries (ADL '98)*, p256.

Bier, E. A., & Perer, A. (2005). Icon abacus and ghost icons. *Proceedings of the First ACM/IEEE-CS joint conference on Digital libraries*, Denver, CO, USA.

Brath, R., Meters, M. & Senior, R. (2005). Visualization for Communication: The Importance of Aesthetic Sizzle. *Proceedings of the Ninth International Conference on Information Visualisation (IV'05)*, London, UK.

Chen, C. (2005). Top 10 Unsolved Information Visualization Problems. In Visualization Viewpoints. *IEEE Computer Graphics and Applications*, July/August, 12-16.

Dang, G., North, C., & Shneiderman, B. (2001). Dynamic queries and brushing on Choropleth Maps. *Proceedings of the Fifth International Conference on Information Visualisation (IV'01)*, London, England, UK.

Grun, C., Gerken, J., Jetter, H., Konig, W., & Reiterer, H. (2005). MedioVis – A user centred library metadata browser. *Proceedings in the Ninth European Conference on Research and Advanced Technology for Digital Libraries (ECDL 2005)*. Vienna, Australia.

Janes, J. W. (1994). Other people's judgements: A comparison of users' and others' judgements of document relevance, topicality, and utility. *Journal of the American Society for Information Science (1986-1998), 43(5)*, 160-171.

Johnson. C. (2004). Top Scientific Visualization research Problems. In Visualization Viewpoints. *IEEE Computer Graphics and Applications*, July/August, 13-17.

Klein, P., Muller, F., Reiterer, H., & Eibl, M. (2002). Visual information retrieval with the

SuperTable + Scatterplot. *Proceedings of the Fifth International Conference on Information Visualisation*, London, England, UK.

Klein, P., Muller, F., Reiterer, H., & Limbach, T. (2003). Metadata visualization with VisMeB. *Proceedings of the Seventh International Conference on Information Visualisation*, Los Alamitos, CA, USA.

Kosara, R., Healey, C.G., Interrante, V., Laidlaw, D.H. & Ware, C. (2003). Thoughts on user studies: Why, how and when. *IEEE Computer Graphics and Applications, 23(4)*, 20-25.

Kumar, V., Furuta, R. & Allen, R. B. (1998). Metadata visualization for digital libraries: Interactive timeline editing and review. *Proceedings of the Third ACM conference on Digital Libraries*, Pittsburgh, Pennsylvania.

Li, Q. & North, C. (2003). Empirical comparison of dynamic query sliders and brushing histograms. *Proceeding of the IEEE Symposium on Information Visualization (InfoVis'03)*, IEEE Computer Society Press, 147-154.

Limbach, T., Muller, F., Klein, P., Reiterer, H., & Eibl, M. (2002). Visualization of metadata using the SuperTable + Scatterplot. *Proceedings of the Internationales Symposium für Informationswissenschaft (ISI) (isiwi)*, Regensburg, Germany.

Microsoft Corporation (2001). *Mircosoft ® Windows*. NSW, Australia: Microsoft Corporation.

Muller, F., Klein, P., Limbach, T., & Reiterer, H. (2003). Visualization and interaction techniques of the visual metadata browser VisMeB. *Proceedings of the third International Conference on Knowledge Management*, Graz, Austria.

NISO (2004). *Understanding metadata*. National Information Standards Organization, NISO Press.

North, C., & Shneiderman, B. (1999a). *Snap-together visualization: Coordinating multiple views to explore information (Technical Report No. #CS-TR-4020)*. University of Maryland Computer Science Department, Maryland, USA.

North, C., & Shneiderman, B. (1999b). *Snap-together visualization: Evaluating coordination usage and construction (Technical Report No. #CS-TR-4075)*. University of Maryland Computer Science Department, Maryland, USA

North, C., & Shneiderman, B. (2000a). Snap-together visualization: Can users construct and operate coordinated visualizations? *International Journal of Human-Computer Studies, 53*, 715-739.

North, C., & Shneiderman, B. (2000b). Snap-together visualization: A user interface for coordinating visualizations via relational schemata. *Proceedings of the Advanced Visual Interfaces (AVI'2000)*, Palermo, Italy.

Nowell, L. T., France, R. K., Hix, D., Heath, L. S., & Fox, E. A. (1996). Visualizing search results: Some alternatives to query-document similarity. *Proceedings of the 19th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'96),* Zurich, Switzerland.

Plaisant, (2004). The challenge of information visualisation evaluation. *Proceedings of the IEEE Conference on Advanced Visual Interfaces (AVI 2004)*, Gallipoli, Italy.

Purchase, H.C., (2000). Effective information visualization: a study of graph drawing aesthetics and algorithms. *Interacting with Computers 13(2)*, 147–162.

Rao, G & Mingay, D. (2001). Report on Usability Testing of Census Bureau's Dynamaps CD-ROM Product. United States Bureau of the Census, Statistical Research Division, Usability Laboratory.

Shneiderman, B. (1994), Dynamic Queries for Visual Information Seeking, *IEEE Software, 11(6)*, 70—77.

Spotfire: Retrieved December 10, 2006, from http://www.sportfire.com

Tweedie, L. (1997). Characterizing interactive externalisations. *Proceedings of CHI'97, ACM Conference of Human Factors in Computing Systems*, Atlanta, 375-382.

Weiss-Lijn, M., McDonnell, J.T., & James, L. (2001). Visualising document content with metadata to facilitate goal-directed search. *Proceedings of the Fifth International Conference on Information Visualisation (IV'01)*, London, England, UK.

Westerman, S.J., Collins, J., Cribbin, T., 2005. Browsing a document collection represented in two- and three-dimensional virtual information space. *International Journal of Human-Computer Studies 62(6)*, 713-736.

Westerman, S.J., Cribbin, T. 2000. Mapping semantic information in virtual space: dimensions, variance, and individual differences. *International Journal of Human-Computer Studies 53*, 765-787.

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA		1. PRIVACY MARKING/CAVEAT (OF DOCUMENT)
2. TITLE The Use of Metadata Visualisation to Assist Information Retrieval		3. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION) Document (U) Title (U) Abstract (U)

4. AUTHOR(S) Agata McCormac, Kathryn Parsons and Marcus Butavicius	5. CORPORATE AUTHOR DSTO Defence Science and Technology Organisation PO Box 1500 Edinburgh South Australia 5111 Australia

6a. DSTO NUMBER DSTO-TR-2057	6b. AR NUMBER AR-014-031	6c. TYPE OF REPORT Technical Report	7. DOCUMENT DATE October 2007

8. FILE NUMBER 2007/1011273/1	9. TASK NUMBER INT 07/020	10. TASK SPONSOR A2C (R5-2-11) PBB (R6-3-07) PB (R6-3-10)	11. NO. OF PAGES 28	12. NO. OF REFERENCES 35

13. URL on the World Wide Web http://www.dsto.defence.gov.au/corporate/reports/DSTO-TR-2057.pdf	14. RELEASE AUTHORITY Chief, Command, Control, Communications and Intelligence Division

15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT

Approved for public release

OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE, PO BOX 1500, EDINBURGH, SA 5111

16. DELIBERATE ANNOUNCEMENT

No Limitations

17. CITATION IN OTHER DOCUMENTS	Yes

18. DSTO RESEARCH LIBRARY THESAURUS http://web-vic.dsto.defence.gov.au/workareas/library/resources/dsto_thesaurus.htm

cognitive psychology
human computer interaction
metadata
visualisation
information retrieval

19. ABSTRACT
Metadata is descriptive information about data, which can be used to manage, locate or retrieve information. Although tabular presentations of metadata can be extremely useful, the exploitation of such information may be improved by visualisation. There are a number of information visualisation interfaces available, and many of these utilise metadata. However, on the whole, these approaches have not been objectively evaluated, and there is little information about their validity and reliability. This study analyses some of these techniques, highlighting their strengths and weaknesses, and the areas where further research is required. A study is proposed, in which a simple interface will empirically assess the efficacy of a metadata visualisation technique.